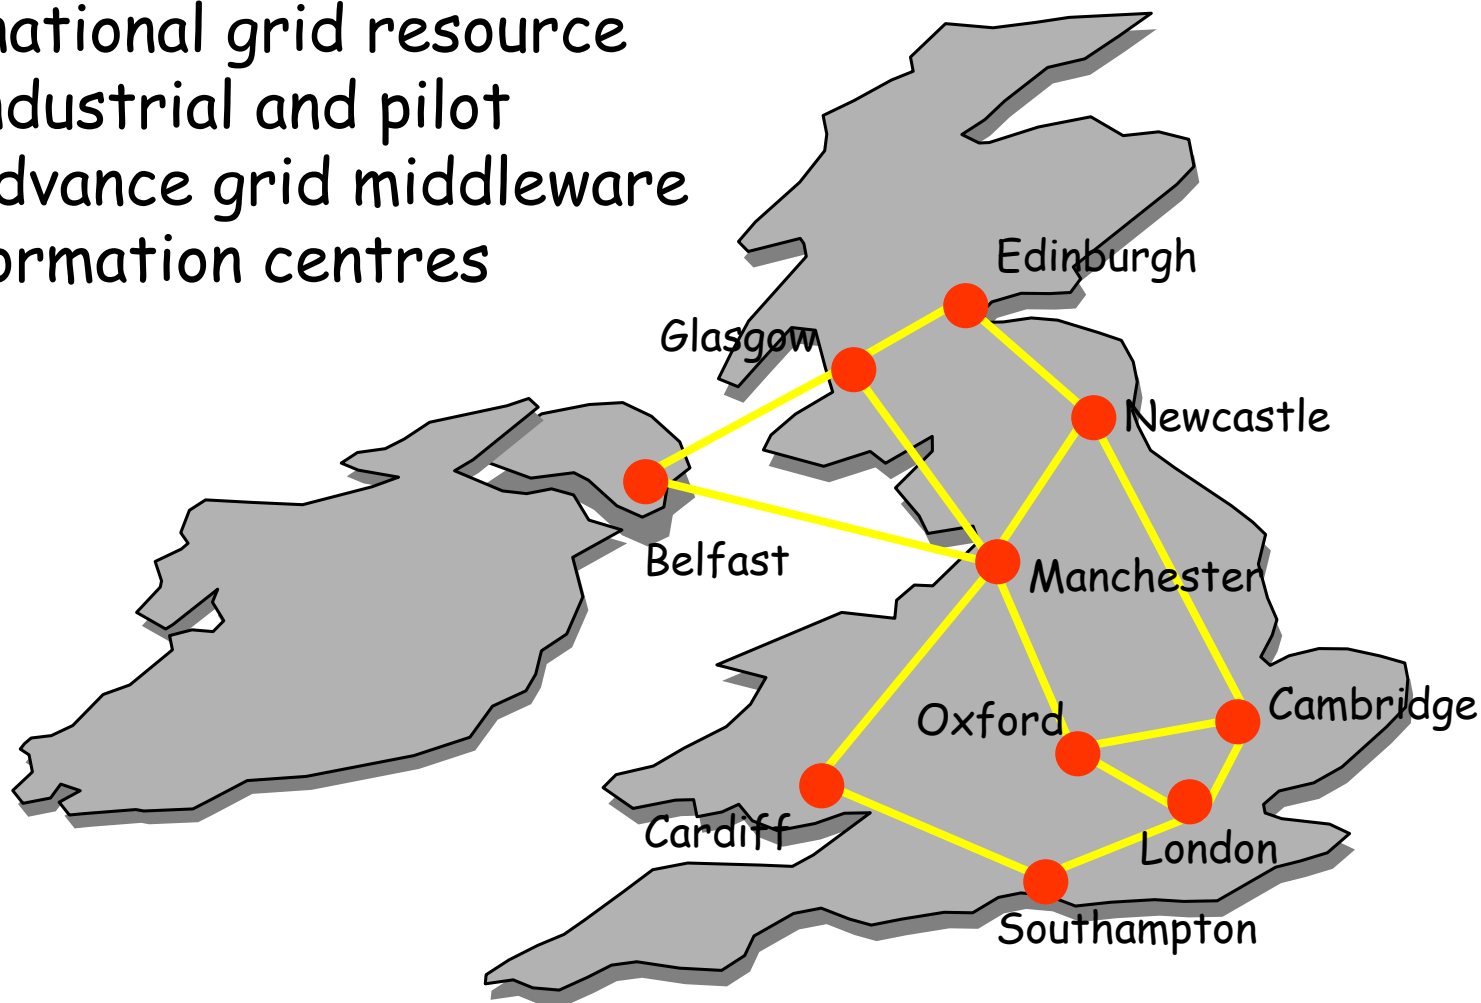


Grid Performance Issues in BeSC e-Science Projects

Ron Perrott & Terry Harmer

- Provide a national grid resource
- Through industrial and pilot projects advance grid middleware
- Act as information centres



- Academic Project
 - GT2 based
 - Large scale atomic physics calculation
- Commercial Projects
 - GT3/web services based
 1. Broadcast Media
 2. High performance data cleansing and mining

- **Academic Project**
 - **Large scale atomic physics calculation**
 - Research tool developing a grid solution for a legacy high-performance computer system
 - **Comparison with supercomputer implementations**
- **Commercial Projects**
 - **Broadcast Media**
 - Prototyping large scale data movements
 - Timings based on broadcast infrastructure scenarios
 - **High performance data cleansing and mining**
 - Prototyping (and using commercially) cleansing and mining
 - Mining/cleansing technology based on utility computing and remote service provision


2DRMP-G

The 2-d R-matrix Propagator Application

- What is 2DRM used for?
 - Calculation of atomic collision data
 - Data is useful in analysis of physical phenomena
 - Laser physics, plasma physics, atmospheric physics, astronomy
 - Test ground for computational approaches
- Why?
 - Relatively little accurate cross-section data is known for many of the processes that are modelled
 - Accurate simulation of electron impact and excitation is difficult
 - Recent theoretical and computational advances enable the accurate study of electron collision for H-like atoms at intermediate energies

- Development started in Autumn 2002
 - GT2
- Deployed on UK grid January 2003
- Operational as a physics tool since March 2003
- Running almost every day since March 2003
 - As a computational tool and a robustness tool for UK L2G

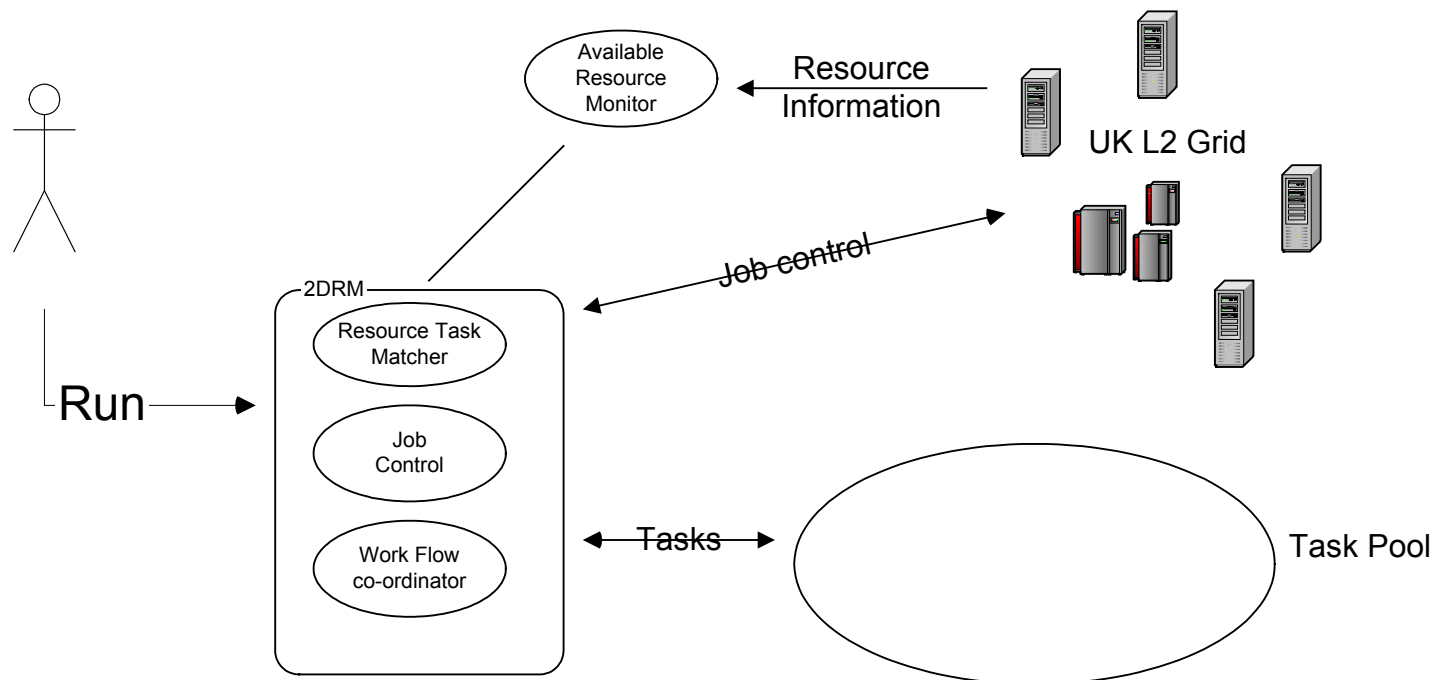
The Solution Structure



Collection of stages organised as conceptual blocks A to D


- Stages 1 & 2 construct atomic basis functions
- Stage 3 constructs sub-region Hamiltonian matrices
- Stage 4 diagonalises the sub-region matrices
- Stage 5 constructs surface amplitudes
- Stage 6 propagates local R-matrices across sub-regions
- Stage 7 is a 1-D propagation step

2DRM-G Solution Architecture

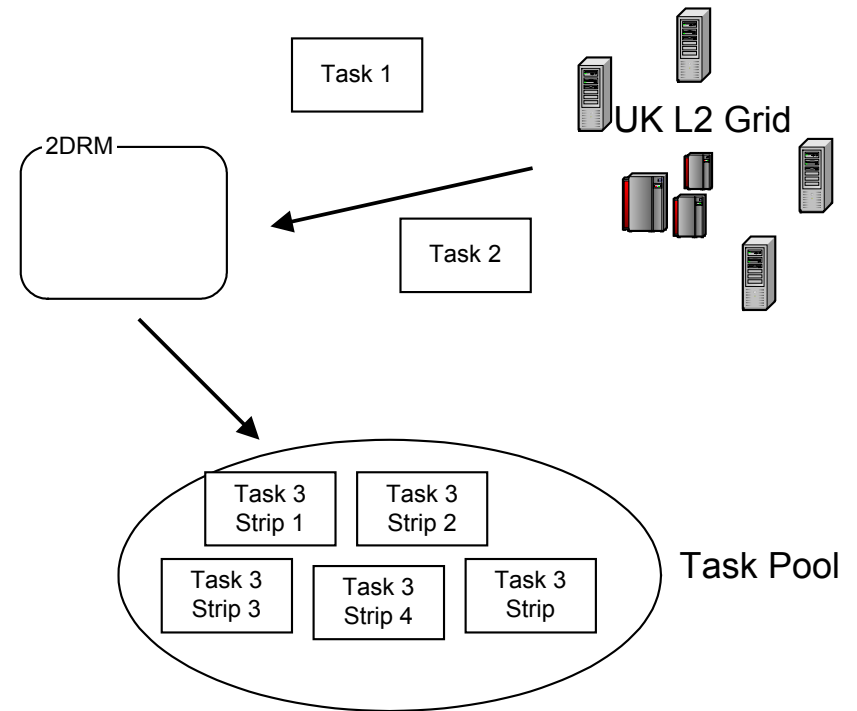


- A client that enables job monitoring.
- A resource list (from UK MDS), pool of tasks, resource-task matcher, workflow co-ordinator and job controller.

Solution Operation




Tasks 1 and 2 are matched with available resources




Completion of Tasks 1 and 2 schedules all of Task 3

Execution Time: Averages Over Last 10 Months

Time / Complete Run



Newrd & Diag Average Times



Results Turn Around Time

...the controversial bit...

Machine	Time
HPCX	50 secs Queue time ~1 hour
L2G	About 1 hour

GridCast


Using the Grid in Broadcast Media Infrastructures

British Broadcast Corporation

British Telecom (NI)

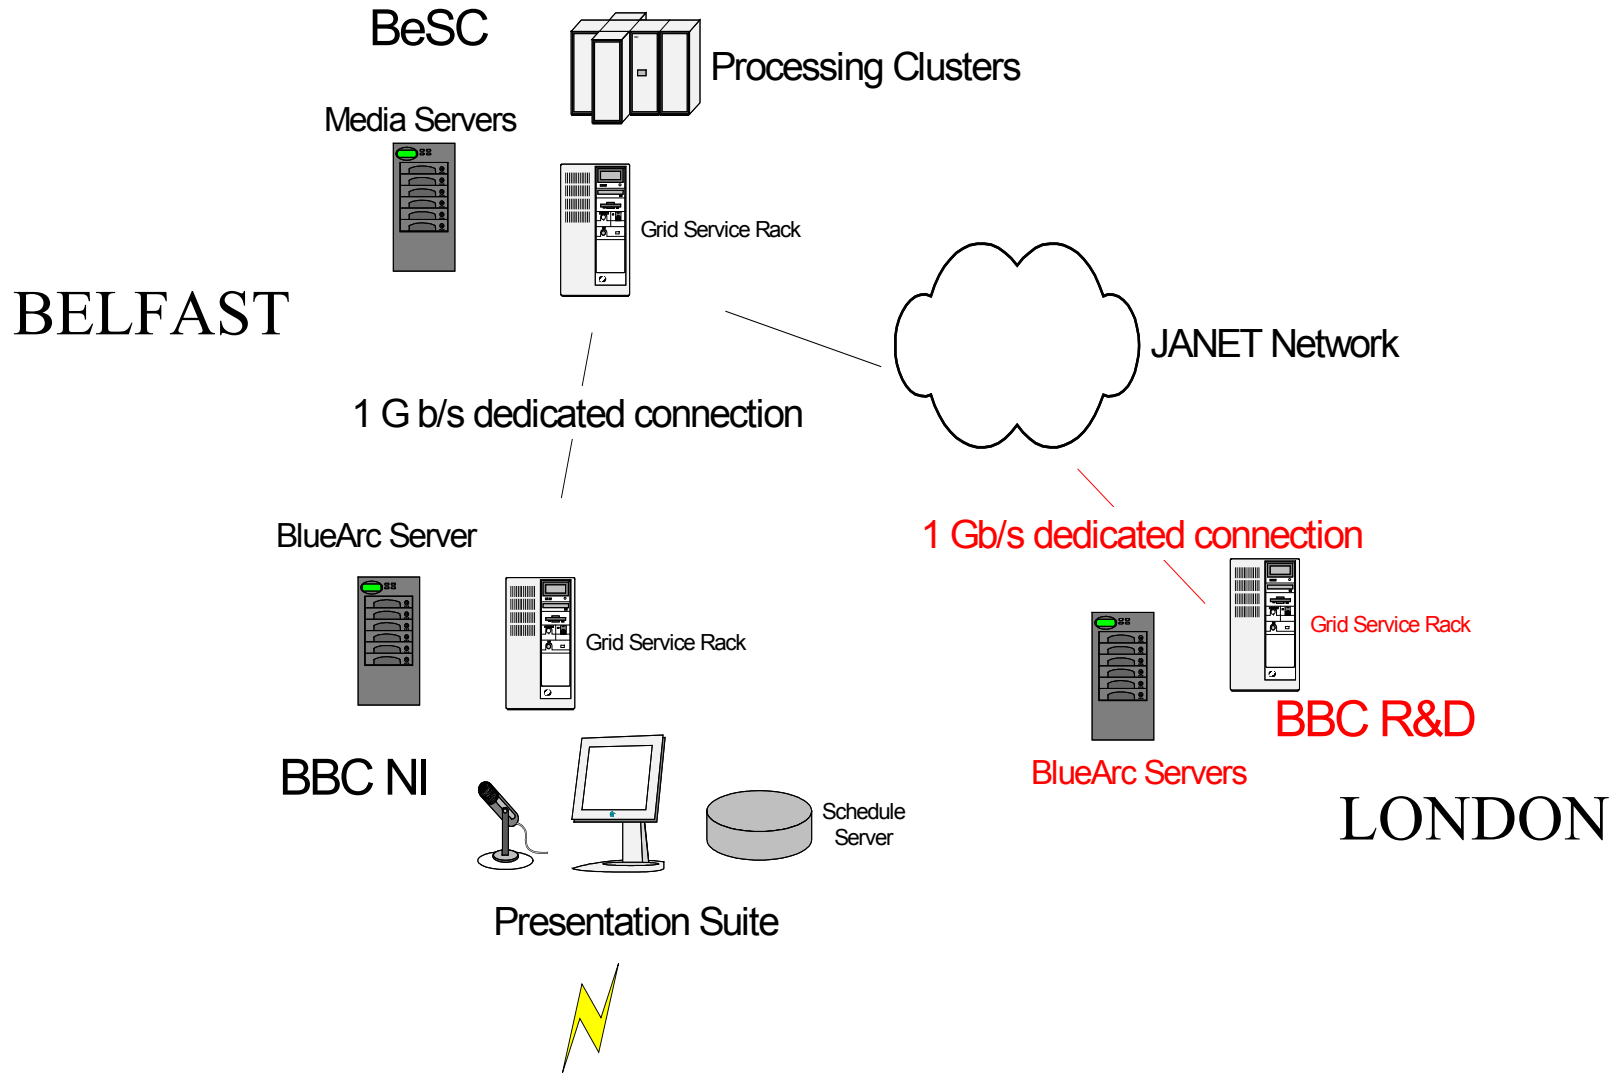
- To develop a baseline media grid to support a broadcaster
 - *Manage distributed collections of stored media*
 - Prototype security and access mechanisms
 - Integrate processing and technical resources
 - Integrate with media standards and hardware
- To analyse *Quality of Service* issues
 - Analyse remote content distribution infrastructures
 - Measure QoS issues in distributed media collections
 - Analyse remote service provision
 - *To analyse reactivity, reliability and resilience issues in a grid-based broadcast infrastructure*

The Grid Scenario: The BBC Nations BBC NI, BBC Scotland and BBC Wales



- **BBC Nations** provide customised services in each nation
- **Television** programmes are distributed to BBC Nations from BBC Network (London) using dedicated leased ATM circuits.


- Project Started in April 2003
 - GT3/Web services based
 - Starting as a GT3 Alpha system in March 2003!
- Infrastructure is in place and the QoS testing is ongoing
 - A small part of this is reported here




- Reflects broadcast infrastructure pattern of use.
 - Intense short periods of activity
 - Reflects usage to co-ordinate broadcasting of content
 - At programme broadcast junctions

Domain Creation Average Timings


Pie chart showing break-down of service activity timings on a domain



Registration Timings over time for the Public Content Service




Registration Timings over time for the Replica Service



- Invoking a service securely takes several seconds
- Not really a problem (currently) in GridCast because we are interested in managing very large media content files
 - Currently 25Gbyte per hour to 180Gbyte per hour
 - Reactivity for non content-based services will be an issue.

GEDDM

Grid Enabled Distributed Data Mining



datactics

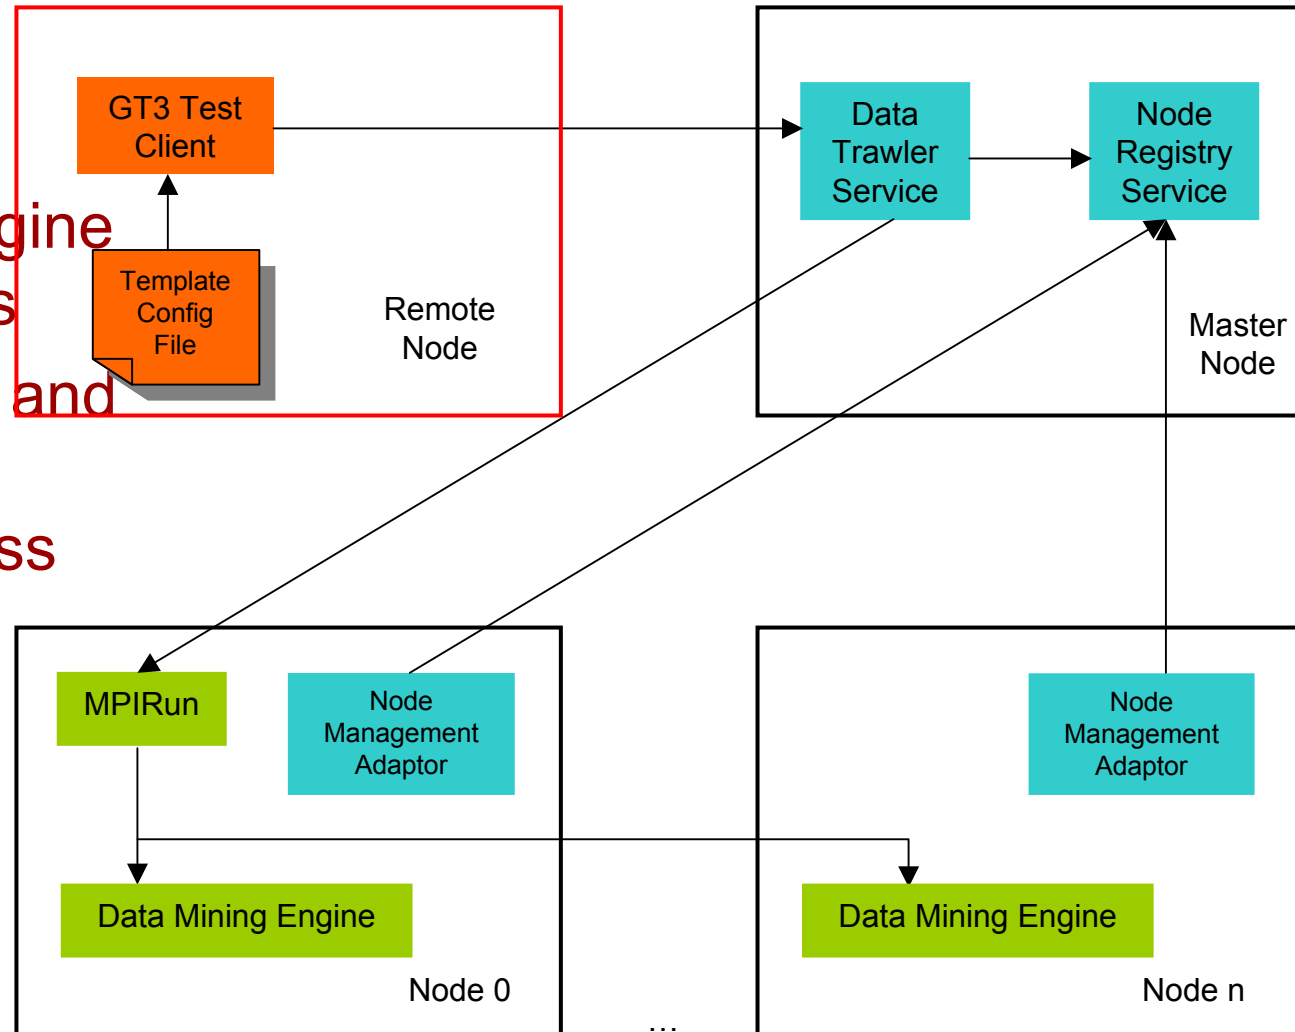
Innovations in Data Matching

- Distributed Data Mining and Conversion of Unstructured data
 - Email, pdf, weblogs, RDBMS, Word, files ...
- Specialising in grid enabled “data-centric” matching across multiple sectors
- *Computationally intensive - need to compare every record with every other record*
- Improve data quality by applying fuzzy matching and parallel processing to achieve greater depth and accuracy
 - Scenarios based on
 - Forensic accounting, Banking, anti-terrorist, crime
- Data mining software being used in the real world



- Mining engine has been in commercial use since 1999
- GEDDM project in development since September 2003
 - GT3/web services based
 - Interested in utility computing and remote mining service provisioning.

- Use Grid Technology
- Expose core engine as Grid Services
- Secure, reliable and scaleable High bandwidth access

Mining Engine



- Java client connecting to grid services
- Secure connection uses GT3 Message Level Security



- Core Data Mining Engine
 - 1 Million records – 3 seconds
(2 seconds – initialisation, 1 sec matching)
- Grid Services
 - GT3 Java, x509 etc
 - Client to DT service
 - Service access time **4 seconds**
- ***Service access time is an issue....***